H. QUALIFICATIONS AND BIOGRAPHY

Robert G. Harris is an Associate Professor in the Walter A. Haas School of Business, University of California, Berkeley, and a Principal in the Law and Economics Consulting Group. He earned a Doctor of Philosophy degree in Economics from the University of California, Berkeley. At Berkeley, he teaches undergraduate, MBA and Ph.D. courses in Business & Public Policy; Economics for Managerial Decisions; Antitrust and Economic Regulation; and Telecommunications Economics, Policies and Strategies (http://haas.berkeley.edu/~wba212/). He has published several dozen articles and papers analyzing the effects of public policies on industry performance in telecommunications industries. He is Co-Director of the Consortium for Research in Telecommunications Policy, a collaborative program of the University of California at Berkeley, the University of Chicago, the University of Michigan and Northwestern University.

He has testified before committees of the U.S. Senate and the U.S. House of Representatives on antitrust, transportation and telecommunications legislation; to the Federal Communications Commission on spectrum auction rules, transfer of cellular licenses, price cap regulation, deployment of broadband networks, dominant carrier regulation and cable rate regulation; to the public utility commissions of fourteen states on rate design, price cap regulation and local competition policy; to the Canadian Radio-Television and Telecommunications Commission and the Mexican Secretariat of Communications and Transportation on interconnection and access pricing; and to numerous U.S. District Courts in antitrust, intellectual property and other business litigation. He has served as a consultant to the U.S. General Accounting Office, U.S. Office of Technology Assessment, U.S. Department of Justice, U.S. Department of Transportation, Interstate Commerce Commission, California Attorney General and California Department of Consumer Affairs. He has also consulted to numerous telecommunications carriers and equipment manufacturers.

ATTACHMENT B

The Internet: A Case Study of Interconnection Pricing in a Nonregulated Competitive Market

The Federal Communications Commission (FCC) was recently told that Internet service providers exchange traffic with each other for free:

> Commercial Internet service providers agreed that interchange of traffic among them was of mutual benefit and that each should accept traffic from the other without settlements payments or interconnection charges. The CIX members therefore agreed to exchange traffic on a "sender keep all" basis in which each provider charges its own customers for originating traffic and agrees to terminate traffic for other providers without charge.
>
> The Internet example suggests that "sender keep all" interconnection arrangements are likely to develop in competitive communications markets as the compensation method for mutually beneficial interconnection arrangements.[1]

This paper demonstrates that the assertion that Internet service providers exchange traffic for free is erroneous. In fact, the Internet generally follows a model of asymmetrical compensation arrangements in which smaller networks (those occupying a lower level position on the Internet hierarchy) pay larger (or higher level) networks for the privilege of interconnection. Thus, to the extent that the Internet is the "best existing example of interconnection under competitive conditions without regulation,"[2] the Internet example suggests that freely com-

[1] Gerald W. Brock, The Economics of Interconnection, at i-ii (April 1995). *See also* Gerald W. Brock, Price Structure Issues in Interconnection Fees, at 1-2 (March 30, 1995). The FCC cites these and other Brock papers in its recent Notice of Proposed Rulemaking, Interconnection Between Local Exchange Carriers and Commercial Mobile Radio Service Providers, CC Docket No. 95-185, FCC 95-505 (Jan. 11, 1996).

[2] *See* Brock, The Economics of Interconnection, at i.

petitive markets will impose asymmetrical interconnection charges among carriers of different sizes.

Internet networks have historically used fixed capacity charges in their compensation arrangements with each other, with the size of the fee based on the capacity of the connecting facility. A customer (*e.g.,* LAN) or small network pays this flat-rated fee in exchange for unlimited access up to the physical maximum throughput of the particular connection (bandwidth). That fixed connection fee approach has been used largely because of the transport technology used within the Internet (packet-switching).

While the connection fees have historically been capacity based, the exponential growth in traffic and the new types of traffic that will be transported in the future have resulted in an active discussion within the Internet community about whether the Internet should abandon capacity charges or supplement them with usage-based pricing to help control the burgeoning congestion on the Internet. While care must be exercised in comparing the Internet and the Public Switched Telephone Network (PSTN) because of the significant differences in technologies and regulation, the Internet example further suggests that some usage-based pricing may be necessary in addition to capacity-based pricing in the PSTN — if only to avoid in the PSTN the "traffic jams" now being experienced in the Internet and to provide the capital to expand the PSTN so it continues to have the capacity needed to handle future telecommunications needs.[3]

This paper is divided into five sections. The first section provides a working description of the Internet, and section two describes its growth and evolution. The third section notes several differences between the Internet and the PSTN, differences which have important implications for pricing and for the effi-

[3] "The Interminablenet: Why is the Internet so slow? And what can be done about it?", The Economist at 70 (Feb. 3, 1996).

cient use of network resources. Section four examines the current pricing and economics of the Internet, and the final section discusses some of the challenges the Internet faces in the near future as it continues to evolve.

I. The Internet Generally: A Working Description

While the Internet has become the darling of the mainstream media, its actual workings remain shrouded in mystery. The Internet has no defined structure and is best described as a network of networks.[4] It is a loose grouping of interconnecting computer networks that use different languages, yet are capable of communicating using a common protocol. That protocol, the "Transmission Control Program/Internetworking Protocol" (TCP/IP), is the glue that holds the Internet together.[5]

The Internet's most distinctive characteristic is that it is virtually devoid of regulation. There is no centralized Internet governing authority; no state, federal or international agency regulates the Internet. Instead, the Internet manages to thrive through largely bilateral contractual arrangements and through the efforts of user groups that set standards and allocate network addresses.[6]

---

[4] Historically, the most accepted formal definition of the Internet has been the collection of interconnected data networks that run the TCP/IP protocol.

[5] Also critical to the Internet's success has been the use of a common, user-friendly addressing scheme, known as the Domain Name System. The DNS, first developed in 1986, allows users to use easily remembered addresses such as username@companyname.com instead of their Internet Protocol (IP) address, such as 192.135.191.128. In 1993 the National Science Foundation (NSF) gave Network Solutions, Inc. a contract to manage registration of domain names and IP numbers. The explosion in Internet usage resulted in a concomitant explosion in new name requests. During the eight-month period between December 1994 and July 1995, registration requests increased from 3,500 monthly to over 15,000 monthly. In September 1995, the NSF authorized Network Solutions to begin charging $50 annually per domain name so that the cost of administering the Domain Name System would be shifted from taxpayers to users. *See* George Lawton, "InterNIC Begins Charging for Domain Names," 12 Digital News & Review, No. 18, p. 7 (Oct. 9, 1995); Kara Swisher, "The Frenzy Over the Internet's Fee Enterprise," The Washington Post, Financial, p. F10 (Oct. 16, 1995); "Dealing with the Name Game Dilemma," 1 Internet Week, No. 29 (Oct. 23, 1995).

[6] The coordination and standardization of Internet protocols is overseen by the Internet Activities Board (IAB) and its two task forces: the Internet Engineering Task Force (IETF), which estab-

Continued on Next Page

While the Internet has not been subject to regulation, the federal government has nevertheless played a pivotal role in its creation, development, and growth. As discussed below, the federal government has spent millions of dollars subsidizing the original Internet backbone network, funding the development of regional networks and, more recently, promoting a more de-centralized network architecture through the installation of new regional network access points (NAPs), or hubs, where Internet traffic is exchanged. This government funding has been instrumental in the Internet's development and has impacted the evolution of the commercial Internet as we know it today — including the early compensation arrangements between Internet carriers.

The Internet is in many ways a virtual network, for it has few physical facilities of its own. Instead, it piggy-backs on the facilities of private and public computer networks, interconnecting them with dedicated lines leased from local telephone companies, competitive access providers, and interexchange carriers (IXCs).

The Internet is traditionally envisioned as having three hierarchical levels: local networks; regional networks; and national transit or backbone networks. At the bottom of the pyramid are Local Area Networks (LANs) and local Internet service providers. The first local networks were LANs, short-distance data communications networks which link personal computers (PCs) and enable them to share resources. LANs often have a central server holding network resources. If the LAN is connected to the Internet, the server connects to a regional network (or another Internet access provider) through a specialized computer known as a router. In this model, the LAN server is a "node" or "host" on the Internet.

Also proliferating in recent years are local Internet service providers (ISPs) which provide dial-up (and dedicated) Internet access to residents and busi-

---

lishes Internet protocol and architecture standards; and the Internet Research Task Force (IRTF), which develops protocols and technology for the future.

nesses within the local ISP's community.[7] Local ISPs, or resellers as they are commonly known, lease a private line (*e.g.*, DS-0, DS-1, or fast packet service (*e.g.*, frame relay, SMDS) from their office to a regional network (or another Internet access provider) so their customers can connect to the global Internet.

A regional network usually serves a specific geographic area often encompassing one or more states. A regional network uses leased lines to interconnect LANs and local Internet service providers (ISPs), and linking them to one or more national Internet backbone providers. The regional networks connect directly to the backbones, which provide cross-country carriage along with connections to other regional networks and other national and international backbones. In recent years, some regional networks also provide Internet access and services directly to businesses and residential customers.

Until recently, national backbone services were provided by such unfamiliar firms like AlterNet/UUNet, ANS, and PSI, which leased private lines from IXCs. More recently, certain IXCs (*e.g.*, SprintLink, InternetMCI) have begun using their own facilities to provide a backbone Internet service.

To give an example of an Internet transmission, students at MIT connect to MITNet, a high-speed optical fiber LAN linking computers all over MIT's campus. One of those computers operates as a gateway to NEARNet, a regional network serving the northeastern states. NEARNet, in turn, may connect to a backbone provider such as MCI. Thus, if an MIT student attempts to access resources located at the National Supercomputer Center in San Diego, the student's inquiry is carried over MITNet, then NEARNet, then MCI. At one of the exchange, or network access, points around the country, the student's inquiry (a

[7] For example, Summit County, Colorado (pop. 14,000) has three local ISPs serving the community: ColoradoNet, Colorado SuperNet (also a regional network), and VailNet. It is not financially attractive for Summit County residents to connect to commercial electronic service providers like CompuServe and America Online because their nearest access point is in Denver, which is a toll call. Local ISPs like ColoradoNet allow Summit County residents to access the global Internet using their local telephone service (thereby avoiding toll charges).

packet) is switched to the backbone provider serving the destination host, in this case, the Supercomputer Center. The inquiry is then delivered to the appropriate regional network which delivers the transmission to the Supercomputer Center.

This description, however, presents a somewhat misleading picture of rigid hierarchy. Many regional networks now connect to one another directly, exchanging traffic without using a backbone provider.[8] So long as they adhere to the same TCP/IP protocol, LANs also can be connected with each other or connected directly to a backbone network, bypassing the regional networks. These interconnections between networks at the same hierarchical level have blurred the old hierarchical model of the Internet. It is for this reason that the Internet today is more aptly described as a cloud or a web.[9]

Whether between networks at the same level or networks at different levels, internetwork connections are performed by specialized computers known as routers. These computers perform routing functions for the Internet that are roughly analogous to the role of a switch in the public switched telephone network. For instance, routers determine the best route between any two networks, even if there are numerous networks in between. Routers also provide network management services such as load balancing, traffic statistics, and prioritization.

---

[8] Many Internet networks connect to more than one network (*e.g.*, a regional may connect to two or more backbones) to provide redundancy in the event of network failures in one network. These multiple connections complicate considerably the development of efficient routing tables. *See* "The Interminablenet: Why is the Internet so slow? And what can be done about it?", The Economist at 70 (Feb. 3, 1996)("Each router creates another path in the network, which all the other routers in the next layer of network must keep track of. Updating these internal 'tables' as the network constantly changes can become a router's toughest job.").

[9] *See, e.g.*, Jeffrey K. MacKie-Mason, Economic FAQs About the Internet, "What is the structure of the Internet" (July 11, 1995)("The new 'privatized Internet' in the US is becoming less hierarchical and more interconnected. The separation between the backbone and the regional network layers of the current structure are blurring, as more regionals are connected directly to each other through network access points (NAPs), and traffic passes through a chain of regionals without any backbone transport.").

Unlike the traditional circuit-switched telephone network, however, the Internet is a packet-switched network. Individual transmissions are subdivided into discrete packets using the TCP protocol, sent through the Internet according to the Internet Protocol, or "IP," and reassembled at the terminating computer.[10]

In the first step, the data is broken down into packets of varying size, averaging 200 bytes according to the TCP protocol. Each individual packet carries its own identifier, which enables the computer on the terminating end to reassemble the packets in the correct order, again using the TCP protocol.[11] This packet-specific identification also facilitates error correction — if a packet fails to show up, or arrives in damaged form, the receiving computer instructs the originating computer to re-send the suspect packet.

To ensure that the entire data transmission reaches its destination, every component packet also contains a header containing the address of the recipient computer. Every host computer on the Internet has a unique address defined in accordance with the Internet Protocol.[12] Routers, which sit between the networks, read the address in the header and determine the quickest route to the packet's destination. Because each packet bears the ultimate address, the routers are able to use an opportunistic routing mechanism — as a packet arrives, the router determines which path which is most efficient at that particular time.

---

[10] The TCP is roughly equivalent to Layer 4 of the OSI protocol stack, while the IP corresponds to OSI Layer 3. IP handles network addressing and fragmentation control; TCP ensures reliable delivery of the message. Because IP resides at the network layer and TCP at the transport layer, TCP/IP can be used on a variety of physical networks, such as Ethernet, Token Ring, FDDI, SMDS, and ATM.

[11] As noted, the average packet contains about 200 bytes of data. On top of these 200 bytes, the TCP/IP headers add about 40 bytes. Thus, about 17% of the traffic carried on the Internet is simply header information. *See* Jeffrey K. MacKie-Mason & Hal R. Varian, Some Economics of the Internet, at 4 (Feb. 17, 1994).

[12] Not every user with Internet access is actually on the Internet. In our MIT example, the student's PC is not an Internet host. Instead, the individual PCs are attached to an MIT server, which could be programmed to operate as MIT's Internet host. This host is attached to the regional network through a gateway; the host, which can be accessed by any Internet user throughout the world, is MIT's Internet node.

Because the networks comprising the Internet often are interconnected at numerous points, the most efficient path, in principle, can vary from second to second, and even from packet to packet. Successive packets may travel along vastly different routes to reach the same ultimate destination, where the recipient computer reassembles the packets in their original format.

As presently configured, the Internet does not allow for prioritization of any one packet or stream of packets.[13] Unlike a circuit-switched network, which provides a temporarily dedicated circuit, a packet-switching network mixes all the traffic. For this reason, the Internet is often compared to a pipe: the packets that make up one data transmission are dropped into the pipe along with other packets that make up other data transmissions. There will be no busy signal if users clog the Internet by sending an abnormal amount of data; subsequent packets will simply queue up at the routers and encounter delays. In this regard, today's Internet offers what is known as "best effort" delivery, not guaranteed delivery.

II. History of the Internet

The Internet is a case study of a collection of networks in constant evolution. The Internet of today bears little resemblance to the Internet of even five years ago.

At its conception, the Internet was a computer transport network devoted to the military, after which it evolved to a government-subsidized network for broader non-military research and education. Thereafter, the Internet became a

[13] The next generation of the IP protocol — IP version 6, or IPng as it was initially known — may support such prioritization. Under this new scheme, the flow ID in each packet header would be designed to accommodate interactive applications by aggregating packets that need a common quality of service. Flow ID is an intermediate step between a circuit setup with dedicated capacity and a pure datagram mode of operation, like the Internet uses today. This feature could be used not only to reserve bandwidth for these packets (in circumstances like video where time sensitivity is critical) but also to attach premium pricing to such transactions. *See* Christine Hudgins-Bonafield, "How Will the Internet Grow," CMP Publications, Inc.

non-commercial, federally subsidized backbone with parallel commercial backbones using few interconnection points or hubs. Today, the Internet consists of a complex system of multiple commercial backbones and networks and a growing number of major exchange points. A brief history of this evolution is presented below.

A. ARPANET: 1969-1985

The Internet had its genesis in 1969 as a military network connecting several hundred computers. Established by the Defense Advanced Research Projects Agency (DARPA), the network was designed to meet a perceived problem in the Department of Defense (DoD): how to keep military sites across the country in communication in the event of a nuclear war. Because this packet-switched network was funded by ARPA, the network was generally referred to as ARPANET, and access was limited to the military and DoD contractors.[14]

ARPANET originally used a wide variety of protocols. The TCP/IP (Transmission Control Protocol/Internet Protocol) was developed in the mid-1970s, and by 1983 all nodes on ARPANET were required to use TCP/IP. Also in 1983, the DoD separated the unclassified portions of the Data Defense Network from ARPANET, creating MILNET.

B. NSFNET: 1986-1994

The National Science Foundation (NSF) established an office for networking in 1984, and it soon decided to fund a national high-speed data network linking several supercomputer centers around the United States. NSF's first ef-

[14] *See, e.g.*, Jeffrey K. MacKie-Mason & Hal R. Varian, Economic FAQs About the Internet, at 1 (Aug. 21, 1994); Jeffrey K. MacKie-Mason & Hal R. Varian, Some Economics of the Internet, at 1-2 (Feb. 17, 1994);

forts at a national backbone network, known as NSFNET, was a 56 kbps (DS-0) network deployed in 1986.[15]

In 1987 NSF contracted with Merit Network, Inc., a non-profit corporation formed by several Michigan universities, to build a national network using T1 (1.5 Mbps) links. Merit subcontracted much of the work to IBM and MCI, with IBM supplying the routing equipment and MCI providing the trunk lines. In 1990 Merit, IBM, and MCI formed a non-profit company, Advanced Network Services, Inc. (ANS), which owned and operated the NSFNET backbone on behalf of NSF. Beginning in 1992, the NSFNET was upgraded to a 45 Mbps (DS-3) backbone.[16]

Because the NSFNET backbone was federally funded, the NSF developed an "acceptable use policy," or AUP, which restricted access to non-commercial research and education uses. The AUP was vague, and anyone who agreed to abide by the AUP was assumed to have done so.

As word spread of the resources available over this web of interlinked computers, a variety of state and regional networks emerged to provide access to user organizations in their geographic areas. Most of these networks were non-profit organizations affiliated with major research universities, and many were formed with NSF funding.[17]

But, as the popularity of internetworking increased, a number of commercial organizations formed to provide Internet access. Some of the regional networks also began offering commercial access. Commercial access, however, was severely hampered by the NSFNET's AUP, which still reserved NSFNET's

---

[15] *See, e.g.*, Jeffrey K. MacKie-Mason & Hal R. Varian, Some Economics of the Internet, at 1-2 (Feb. 17, 1994).

[16] *See* "Another Reading of the NREN Legislation," 25 Telecommunications, No. 11, p. 29 (Nov. 1991).

[17] For example, CERFnet, a California-based regional network and one of the CIX founders, was launched in 1989 with a $2.8 million NSF grant. *See* CERFnet Press Release, Business Wire (March 27, 1991).

facilities for research and education uses. This policy, although rarely enforced, prevented commercial access providers from exchanging traffic *via* the NSFNET backbone.

Internet usage exploded in the early 1990s, and the NSF determined that the provision of transport, or transit, services would best be left to commercial backbone providers. It decided to cut all funding for the NSFNET effective April 30, 1995, when the NSFNET ceased to exist.[18] This left commercial backbone providers — the original networks, AlterNet, ANS, and PSInet; and the more recent entrants, SprintLink, InternetMCI — to take up the slack. Local and regional networks wishing to connect to a backbone had to connect to one of these commercial backbones.

C. The First Commercial Networks: 1991-1994

As noted, the NSFNET's "acceptable use policy" prevented commercial access providers from exchanging commercial traffic *via* the NSFNET backbone even though most of them were connected to NSFNET. The growing need for the transport and exchange of "AUP-free" commercial traffic led to two developments in 1991: the formation of ANS CO+RE, and the formation of the Commercial Internet Exchange (CIX).

1. <u>ANS CO+RE</u>. In 1991 Advanced Network Service, which built and managed the NSFNET, created a for-profit subsidiary, ANS CO+RE Systems, Inc., to develop a DS-3-based Internet backbone and market it for commercial use. Using excess capacity on the NSFNET, the "CO+RE service" was designed to provide a public TCP/IP network service which corporate customers could use

[18] Instead, the NSF began funding a new very-high-speed backbone network (vBNS) to interconnect its five supercomputing centers at 155 Mbps. This network does not provide the regional networks with backbone service, but it does have an acceptable use policy that emphasizes developing capabilities for high-definition remote visualization and video transmission. Various other government agencies continue to maintain backbones that are reserved for their own traffic, such as ESNet (the Energy and Science Network) and NSI (NASA Science Internet).

for any purpose for which they were willing to pay.[19] The advantage of ANS CO+RE was that existing providers could use the same routing procedures as their existing NSFNET traffic — which was not the case for traffic using the CIX (discussed below). Another advantage was that ANS was developing a 45-Mbps (DS-3) backbone to be available throughout the U.S.

As noted above, the NSFNET ceased to exist in April 1995 when NSF subsidies ended. ANS continued operations, however. Indeed, six months earlier, in November 1994, America Online spent $35 million to purchase ANS.[20]

2. CIX. Also in 1991, three prominent service providers announced formation of an "AUP-free" interconnection point — the Commercial Internet eXchange (CIX) — through which member organizations could interconnect without having to access the NSFNET backbone. The three founding members were General Atomics, which managed the California Education Research Federation Network (CERFNet), a regional network; and two national backbones: Performance Systems International, Inc. (operating PSInet), and Uunet Technologies, Inc. (operating AlterNet).[21] At the time of this announcement, these three providers claimed they provided "nearly 100 percent" of the commercial TCP/IP internetworking services in the United States.[22]

---

[19] *See, e.g.*, Eric Arnum, "Transformations in Public Data Networking," 21 Business Communications Review, No. 12, p. 43 (Dec. 1991); "NSF Embroiled in Commercialization Debate," 12 Communications Daily, No. 24, p. 4 (Feb. 5, 1992).

[20] *See* Joanne Cummings, "Internet Service Providers to Ride a Familiar Roller Coaster," 25 Business Communications Review, No. 1, p. 67 (Jan. 1995). About this same time, MCI, which had been one of the ANS owners, announced formation of its Internet network and service, Internet-MCI. Id.

[21] The next year, in January 1992, US Sprint Communications announced that it had formed SprintLink, a TCP/IP router-based network, and had joined CIX. *See* Bob Wallace, "US Sprint Plans Wide-area LAN Internet Service," Network World, p. 5 (Jan. 27, 1992). The following year Sprint announced expansion of its SprintLink Internet service to 12 countries. *See* "Sprint Expands Internet Connections to 12 Countries," 17 Online, No. 5 p. 77 (Sept. 1993).

[22] Business Wire (March 27, 1991).

Initially, CIX consisted of three routers, one owned by each member, located in the San Francisco area. Arranged in a triangle, each router was connected with T1 lines operating at 1.5 Mbps.[23] Soon, however, the CIX deployed its own router in Santa Clara to which members could connect using T1 facilities they obtained.

By 1994, the CIX T1 router links were experiencing considerable congestion, and some characterized the CIX router as "a choke point for delays."[24] During the summer of 1994, CIX announced plans to place its router on Pacific Bell's SMDS cloud in an attempt to relieve this congestion.[25]

D. The Proliferation of Interconnection Hubs: 1995-Future

The CIX was the predominant point of interconnection of commercial Internet networks during the three-year period of 1991-1994. Since that time, the number of interconnection hubs has grown considerably, due to two developments: (1) the NSF's decision to sponsor several regional network access points (NAPs); and (2) the decision by MFS Datanet, Inc. to deploy regional Metropolitan Area Ethernets (MAEs).[26]

1. NSF-Funded NAPs. As part of its privatization program, the NSF decided in 1993 to divert its funding from supporting a backbone network to encouraging the development of several new interconnection points — known as

---

[23] *See* Kelly Jackson, "Team Offering Internet Bypass," Communications Week, Computer Networking Section, p. 18 (April 8, 1991).

[24] Ellen Messmer and Joanie Wexler, "Network Overload Hobbles SprintLink," Network World, p. 6 (Sept. 5, 1994).

[25] *See* Joanie Wexler, "Internet Lands SMDS Backbone," Network World, p. 28 (Aug. 29, 1994); ~~Barry Slawter, "Internet Providers Choose SMDS," 5 Internetwork, No. 9, p. 6 (Sept. 1994).~~

[26] Hubs in addition to NAPs and MAEs are being deployed as well. The most recent phenomenon is the installation of "mini" regional hubs or NAPs for the exchange of primarily regional traffic (*e.g.*, so transmissions between two midwest sites do not have to be routed outside the midwest).

A map of the major interconnection points, and the numerous networks that use them is available at CERFNet. *See* U.S. Network Service Providers Interconnections, Draft V.2.0 (Sept. 10, 1995).

Network Access Points (NAPs) — as an alternative to the CIX router.[27] The NSF took this step to prevent the commercial backbones from setting up a hodge-podge of bilateral connection points — potentially creating routing chaos — and to provide connection points for its new OC-3 vBNS network.[28] Like the CIX router, these new NAPs would be "AUP free."

The NSF decided to sponsor and fund three "priority" NAPs, each selected for its geographic location:[29]

- New Jersey. This Pennsauken, New Jersey-based NAP is operated by SprintLink. SprintLink currently offers a Fiber Distributed Data Interface (FDDI) LAN technology (100 Mbps), although it apparently plans to upgrade to an ATM switch.

- Chicago. This NAP, managed by Bellcore and operated by Ameritech Advanced Data Services, uses ATM for connections up to 155 Mbps (OC-3).

- San Francisco. This NAP, also managed by Bellcore and operated by Pacific Bell, also uses ATM technology at connection speeds up to 155 Mbps.

2. MAEs. Metropolitan Area Ethernets (MAEs) is a new form of interexchange points provided by MFS Datanet, Inc. Unlike the CIX, which consisted of a router, a MAE is a distributed Ethernet service spanning a wide geographic

---

[27] *See* Ellen Messmer, "Internet's Future Lies in NSF Solicitation," Network World, p. 32 (May 24, 1993).

[28] *See* note 18 *supra*.

[29] Although these NAPs will initially be managed under the NSFNET program, NSF funding will terminate in 1998. MFS's MAE East, discussed in the next section, is generally known as a "non-priority" NSF NAP.

area. MAEs are being upgraded from a distributed Ethernet to a collocated FDDI ring (100 Mbps).

The first MAE was MAE-East, located in the Washington, D.C. area, which began as an experimental interconnection arrangement developed by AlterNet, PSI, and SprintLink. Since then, MFS has opened MAE West, located in the San Jose, California area, and MAE Chicago. MFS has also announced plans to open MAE Dallas and MAE Los Angeles.

New network access points are also being installed worldwide. In December 1995, for example, Cable & Wireless announced plans to construct eight NAPs in Australia, Bermuda, Germany, Hong Kong, Japan, Sweden, and two in the United Kingdom.[30]

The current architecture moves the national Internet structure from a noncommercial backbone with parallel commercial backbones using a single commercial interconnection exchange point (the CIX) to a more complex system of multiple commercial backbones with several major exchange points. The addition of more exchange points has allowed for the more efficient routing of traffic (*e.g.*, traffic from one east coast site to another no longer needs to be transported to and from the CIX router in California).

Today, there are six commercial national transit or backbone Internet networks — ANS, AGIS/NET99, InternetMCI, PSI, SprintLink, and UUNET (or AlterNet), known as the Club or Six — which connect to most of the NAPs and MAEs.[31] In addition, there are dozens of regional networks and hundreds of local networks — all of which comprise the Internet.

---

[30] *See* Kenneth Hart, "Web Transforms Carriers' Core Network Strategies," Communications Week International, No. 156, p. 1 (Dec. 18, 1995).

[31] See Gordon Cook, "Peering & Transit at the NAPs and the Club of Six," COOK Report Summary (July 1995). It has been reported that eight regional networks have announced they have

Continued on Next Page

E. Internet Usage and Growth

1. Usage. The most frequent use of the Internet is probably e-mail. After that are file transfer (moving data from one computer to another) and remote login (logging into a computer that is running somewhere else on the Internet). In terms of traffic volume, as of December 1994, about 32% of total traffic was file transfer, 16% was World Wide Web (WWW), 11% was netnews, 6% was e-mail, 4% was gopher, with the rest for other uses.[32]

2. Growth. From 1985 to December 1994, the Internet grew from about 200 networks to well over 45,000 networks, and from 1,000 hosts (computers) to over four million. About 1.0 million of these hosts are at educational sites, 1.3 million hosts at commercial sites, and about 385,000 at government/military sites, all in this country. The remaining 1.3 million hosts are located elsewhere in the world.[33]

NSFNET traffic grew from 85 million packets in January 1988 to 86 billion packets in November 1994.[34] It is difficult to obtain more current Internet growth data since the NSFNET was disbanded. All indications are, however, that the Internet's once astonishing growth has exploded to yet new heights.

III. Differences Between the Internet and the Public Switched Telephone Network

At first blush, the Internet and the Public Switched Telephone Network (PSTN) appear to be alike; after all, the Internet transports telecommunications and uses private lines leased from carriers. Nevertheless, there are important

---

formed a new firm, called Corporation for Research and Enterprise Networking (CoREN), which will also build a national backbone which connects to the NAPs and MAEs.

[32] *See* Jeffrey K. MacKie-Mason, Economic FAQs About the Internet, "What do people do on the Internet" (July 11, 1995).

[33] Id. at "How big is the Internet?".

[34] Ibid.

differences between the Internet and the PSTN, differences in technology and regulation. As explained below, these differences have important implications for pricing and the efficient use of network resources.

A. Differences in Technology. The Internet and the PSTN use very different transport technologies. The PSTN uses circuit-switched technology designed to support voice telephony. A telephone user dials a number, and various switches establish a dedicated path between the caller and the called number. This circuit, with a fixed allocation of network resources, stays open, and no other caller can use those resources until the call is terminated. This means, for example, that a long silence between two teenagers uses the same network resources as an active negotiation between two fast-talking lawyers.[35]

The Internet is a network of networks that use connectionless packet-switching communications technology.[36] A packet-switching network uses statistical multiplexing to maximize use of the communications lines. Each circuit is simultaneously shared by numerous users, and no single open connection is maintained for a particular communications session: part of the data may go by one route while the rest may take a vastly different route.[37]

---

[35] Packet service, in contrast, allows for the efficient use of communications lines. Consider a typical interactive terminal session to a remote computer: most of the time the user is thinking. The network is needed only after a send key is struck or when a reply is returned. Holding an open connection would waste most of the capacity of the network link. Instead, the computer waits until after a key is struck, at which point it puts the keystroke information in a packet which is sent across the network. The rest of the time the network links are free to be used for transporting packets from other users.

[36] "Connectionless" means that there is no end-to-end setup for a session; each packet is independently routed to its destination. Some packet-switching networks are "connection-oriented," notably X.25 and frame relay networks. In such networks, a connection is set up before transmission begins, just as in a circuit-switched network. A fixed route is defined, and information necessary to match packets to their session and defined route is stored in memory tables in the routers. Thus, connectionless networks economize on router memory and connection set-up time, while connection-oriented networks economize on routing calculations (which have to be redone for every packet in a connectionless network). *See* Jeffrey K. MacKie-Mason & Hal R. Varian, Economic FAQs About the Internet, at 5 (Aug. 21, 1994).

[37] The postal service is a good metaphor for the technology of the Internet. A sender puts a message into an envelope (packet) and that envelope is routed through a series of postal stations,

Continued on Next Page

The differences between the PSTN's circuit-switching technology and the Internet's packet-switching technology have critical implications for pricing. For example, a one-minute phone call on a circuit-switched network requires one accounting entry in the usage data base. But in a packet network, that one minute-phone call would require approximately 2,500 average-sized packets; complete accounting for all packets would then require a total of 2,500 entries in the recording data base.[38] It is the inherent suitability of packet service to detailed usage accounting that led to the predominance of connection fees in the Internet, as discussed in the next section.

Another pricing problem concerns the granularity of the records. Presumably, accounting detail is most useful when it traces traffic to the user. Certainly, if the purpose of accounting is to charge prices as incentives, those incentives will be most effective if they affect the person actually making the usage decisions. But the network is at best capable of reliably identifying the originating host computer (just as phone networks identify only the phone number that placed the call, not the caller). Another layer of complex authorization and accounting software would be required on the host computer in order to track which user accounts are responsible for which packets.[39]

Another serious problem for almost any Internet usage pricing scheme is how to determine correctly whether the sender or receiver should be billed. With telephone calls it is clear that in most cases the originator of a call should pay. However, in a packet network both "sides" originate their own packets, and in a

---

each station determining where to send the envelope on its next hop. No dedicated pipeline is ~~opened end-to-end and, thus, there is no guarantee that envelopes will arrive in the sequence~~ they were sent, follow exactly the same route to get there, or even arrive. *See* Jeffrey K. MacKie-Mason & Hal R. Varian, Some Economics of the Internet, at 4 (Feb. 17, 1994).

[38] *See* Jeffrey K. MacKie-Mason & Hal R. Varian, Economic FAQs About the Internet, at 11 (Aug. 21, 1994).

[39] *See* ibid.

connectionless network there is no mechanism for identifying party B's packets that were transmitted as responses to a solicitation by party A.[40]

Users who generate traffic ordinarily should pay the costs associated with that traffic. In the PSTN, the call originator pays unless the receiver agrees to accept a collect call or has an 800 number. This arrangement is much more difficult to track with the Internet where an originator can request that the "called party," such as a Web site, send back an enormous volume of information. As the president of the Internet Society has stated:

> [W]ho is the beneficiary? The sender or the receiver? You could build a case that it is both the sender and receiver. You can't tell from just looking at who initiates the traffic flow who is the beneficiary and who should be the billable party.[41]

B. Differences in Regulation. Another critical difference between the Internet and the PSTN is regulation. Carrier-of-last-resort and universal service obligations require LECs to serve all users at geographically averaged rates regardless of the cost of service. Due to legislative and regulatory directive, certain classes of customers and services (toll, access, business) are usually obligated to cross-subsidize other classes (residential). Furthermore, LECs are required to provide a certain level of service; everyone expects to receive dial tone immediately when they want to make a call, regardless of the day (*e.g.,* Mother's Day) or time of day (*e.g.,* 10 a.m.). LECs are also required to provide enough stand-by capacity to immediately serve new residences and businesses, and to add second lines on demand.

---

[40] Id. at 14.

[41] "Vint Cerf to COOK Report: Discussion Needed of Benefits Derived from Backbone Resources: Fair Compensation to Backbone Providers Must Be Ensured," The COOK Report on the Internet, at 7 (Sept. 1995).

Internet service providers are not subject to any of these regulatory obligations. The Internet "is free of some of the extraordinary requirements (e.g., cross-subsidies, as for support of universal access) found in telephony."[42] ISPs are free to discriminate among their customers and do in fact discriminate between end users and resellers, for example. In addition, Internet providers have no service quality obligations (*e.g.,* capacity available upon demand). The current Internet offers a single service quality: "best efforts packet service."[43] Packets are transported on a first-come, first-served basis with no guarantee of success. Some packets may experience severe delays, while others may be dropped and never arrive.[44]

In summary, while there are some similarities between the Internet and the PSTN, care must be exercised before one concludes that the Internet is "a useful model for interconnection of competing local exchange telecommunications networks."[45]

IV. The Economics and Pricing of Internet Interconnection

Perhaps the most prevalent and persistent misconception about the Internet is that it is "free" — which some take to mean, it is costless.[46] This misconception is fueled in part by the fact that the Internet was once free, funded by taxpayer subsidies. This misconception is furthered by the publicity received by

---

[42] Marjory S. Blumenthal, "Realizing the Information Future: Technology, Economics, and the Open Data Network," in Gerald W. Brock, ed., Toward a Competitive Telecommunications Industry: Selected Papers from the 1994 Telecommunications Policy Research Conference, 275, 284 (Lawrence Erlbaum Associates, 1995).

[43] *See* Jeffrey K. MacKie-Mason & Hal R. Varian, Economic FAQs About the Internet, at 10 (Aug. 21, 1994).

[44] Ibid.

[45] Gerald W. Brock, Price Structure Issues in Interconnection Fees, at 3 (March 30, 1995).

[46] As William Schrader, president of PSI, a CIX founder, stated, Internet "[c]osts are hidden. Some universities have five or six people managing their Internet network, and that's five or six salaries right there." *See* Eric Arnum, "Transformations in Public Data Networking," 21 Business Communications Review, No. 12, p. 43 (Dec. 1991).

the "settlements free" policy adopted early on by the CIX founders. It is also fueled by the fact that many Internet end users, generally those served by campus LANs, often are not charged directly for accessing the Internet.[47]

Of course, the Internet is neither free nor costless. While end users such as our MIT student, for instance, may pay no specific Internet fee, the costs of the MIT LAN (installation and operations) are paid for by MIT. In addition, the cost of connecting the LAN to the Internet is borne by MIT, which pays a flat-fee to NEARNet — the regional provider — for dedicated Internet connection. A medium-sized university typically pays about $60,000 annually in connection fees; these costs are recovered from students, although generally buried in tuition or general fees.

For the most part, the Internet uses asymmetrical compensation arrangements, with money flowing upward, from end users to local and regional networks to backbone networks. "Money flows upwards: Each level pays the next for connectivity and, occasionally, usage."[48]

Nearly all Internet networks use the same two-part interconnection pricing structure for interconnecting with a higher level network in the Internet hierarchy:

1. The network wanting connection is responsible for paying for the connection facility (*e.g.*, a private line leased from a common carrier) connecting the customer's site and the Internet provider's "point of presence" (POP);[49] and

---

[47] Even public users get the sense that each additional piece of e-mail sent to Lithuania or each file retrieved from Denmark is costless, when they pay a flat monthly fee for access.

[48] *See* Kenneth Hart, "Internet Providers Want Body to Manage Growth," Communications Week International (Sept. 1, 1995).

[49] *See, e.g.*, Padmanabhan Srinagesh, "Internet Cost Structures and Interconnection Agreements," Gerald W. Brock, editor, Toward a Competitive Telecommunication Industry: Selected Papers from the 1994 Telecommunications Policy Research Conference, 251, 252 (1995)("[T]he access link is not usually considered a service offered by the ISP.").

2. The network wanting connection pays a fixed fee in exchange for unlimited access up to the physical maximum throughput of the particular connection (bandwidth) — called a connection fee or connection pricing.[50] For example, a local network might pay a regional network X dollars for a DS-0 connection, three times X for a DS-1 connection, and 10 times X for a DS-3 connection.

It is also becoming increasingly common for networks to charge resellers a higher (*e.g.,* three times) wholesale connection fee than that charged to non-resellers.[51] Because resellers, which aggregate traffic from many end users, often generate more traffic than individual direct end users, this multi-tier pricing clearly is a step toward usage-based pricing.

A similar pricing arrangement is generally used for connection to the switching hubs (*e.g.,* NAPs, MAEs). A network wanting to connect to a NAP or a MAE must (1) provide its own connection facility, and (2) pay the NAP or MAE a flat connection fee based on the bandwidth of the connection facility.

A connection to a NAP or a MAE does not, however, guarantee that other networks connected to the hub will "peer" (or interconnect) with a new network.

---

[50] *See* Jeffrey K. MacKie-Mason & Hal R. Varian, Economic FAQs About the Internet, at 8 (Aug. 21, 1994). *See also* Graham Finnie, "Internet Expansion: the Price of Success," Communications Week, at 37 (Oct. 10, 1994)("Most Internet providers do not charge users by traffic volume, or the number of packets sent, but for a contracted, leased bandwidth, generally the capacity of the access circuit to the provider's point of presence.").

In addition to a flat connection fee (monthly or annual), networks often charge a fairly sizable initial connection fee and some impose a separate charge for CPE (*e.g.,* a router to serve as a gateway between the customer's network and the Internet provider's network). *See* Jeffrey K. MacKie-Mason & Hal R. Varian, Economic FAQs About the Internet, at 8 (Aug. 21, 1994). *See also* Padmanabhan Srinagesh, "Internet Cost Structures and Interconnection Agreements," Gerald W. Brock, editor, Toward a Competitive Telecommunication Industry: Selected Papers from the 1994 Telecommunications Policy Research Conference, 251, 254-55 (1995).

[51] *See, e.g.,* Padmanabhan Srinagesh, "Internet Cost Structures and Interconnection Agreements," Gerald W. Brock, editor, Toward a Competitive Telecommunication Industry: Selected Papers from the 1994 Telecommunications Policy Research Conference, 251, 266 (1995). Some backbones also require resellers to use a more complex addressing scheme and routing protocol than required of end users. Id.

A new network must first generally enter into "peering" agreements with all other connected Internet providers so traffic can flow unhindered.[52] As recently explained by the president of MCSNet, a Chicago-based regional network:

> Of course, there is no guarantee that any of these entities will actually peer with you. That's an individual negotiation issue and will need to be worked out in each case. . . . [C]onnecting to a MAE by itself gets you nothing other than a raw pipe. You must also negotiate transit and/or peering agreements, and some of these may involve additional cost depending on who you're talking to and their terms and conditions.[53]

As this MCSNet official further explained, the "typical sticking point would be if you are a dial-up only provider (*i.e.*, you provide no tangible infrastructure that the other party's customers would want to talk to). In that case, there is no reason for a firm to peer with you. For this reason, most peering agreements specify that certain conditions must be present before a firm will peer — some technical, some business-oriented."[54] For example, the transport of another's traffic (transit) "almost always does involve a fee, you are using someone else's infra-

---

[52] SprintLink apparently takes a slightly different approach by requiring connected networks to exchange traffic without charge for the first year:

> The parties agree not to charge the other party for interconnection-related matters, including charges based on traffic volume, commonly called 'settlements', for the initial twelve (12) months of this Agreement. During the initial 12 months of this Agreement, no fees will be charged between or among the parties to pay for digitized information traffic exchanged through the interconnection. The parties agree to work toward a mutually agreeable formula for such fees after this initial period, but such fees will be mutually agreed upon by the parties. Sample Bilateral Interconnection Agreement (current Internet version).

Of course, Sprint charges new networks its connection fees so its additional transport/routing ~~costs are covered~~.

[53] Harold A. Driscoll, "MAE Chicago - An Interview with Karl Denninger," 11 Hard Copy, No. 12, pp. 42-45 (Dec. 1995).

[54] Ibid. A leading Internet observer has stated that the current practice among the large backbones "is not to make unilateral cost free peering arrangements" with other networks. Gordon Cook, "Peering & Transit at the NAPs and the Club of Six," COOK Report Summary (July 1995).

structure and not giving anything back to them for the pleasure (other than your load)."[55]

Simple connection pricing still dominates the market, but a number of variations have emerged. The most notable is "committed information rate" pricing. In this scheme, an organization is charged a two-part fee. One part is based on the bandwidth of the connection, which is the maximum *feasible* flow rate; the second part is based on the maximum *guaranteed* flow to the customer. The network provider installs sufficient capacity to transport simultaneously the committed rate for all of its customers, and installs flow regulators on each connection. When some customers operate below that rate, the excess network capacity is available on a first-come, first-served basis for the other customers.[56]

A. The Economics of the Internet Interconnection

It is understandable that Internet networks have used fixed connection fees as opposed to usage based charges. Internet networks were driven to fixed fees because of the way costs are incurred on the Internet and because of the difficulty of measuring usage in packet-switched networks.

Most of the short-run costs of providing an Internet backbone are fixed.[57] Two components dominate the costs of providing a backbone network: communications lines and routers. Lease payments for lines and routers accounted for nearly 80% of the 1992 NSFNET costs. An additional 7% of the NSFNET

---

[55] Harold A. Driscoll, "MAE Chicago - An Interview with Karl Denninger," 11 Hard Copy, No. 12, pp. 42-45 (Dec. 1995). *See also* Kenneth Hart, "Internet Providers Want Body to Manage Growth," Communications Week International (Sept. 1, 1995).

[56] *See* Jeffrey K. MacKie-Mason, Economic FAQs About the Internet, "What types of pricing schemes are used" (July 11, 1995). This type of pricing is more common in private networks than in the Internet because a TCP/IP flow rate can be guaranteed only network by network, greatly limiting its value unless a large number of the 20,000 Internet networks coordinate on offering this type of guarantee. Ibid.

[57] Of course, Internet service providers incur many costs not incurred by network providers, including sales, billing, and customer relations costs.